Efficient and Robust Machine Learning for Net-Load Forecasting in Low-Voltage Distribution Systems

Anthony Faustine

Supervisor: Dr. Lucas Pereira · Co-supervisor: Prof. Nuno J. Nunes

sambaiga@gmail.com sambaiga/twiga-thesis

TÉCNICO LISBOA

LARSyS Laboratory of Robotics and Engineering Systems

INTERACTIVE TECHNOLOGIES INSTITUTE

FEELab

pypi.org/project/twiga Forecasting library

pypi.org/project/popo-flow Feeder simulation

Outline: four questions

OUTLINE

The problem

Joanna operates an LV substation

Challenges and the goal

SLIDE 1

RQ1 · Evaluation

Can I trust how models are compared?

A five-step evaluation protocol

CH. 3 SLIDES 2–3

RQ2 · Cost

Can I afford an accurate model at every substation?

Lightweight MLPF and MLPGAM

CH. 4, 6

SLIDES 4–6

RQ3 · Coverage

Does a 90% interval achieve 90% coverage?

FPQR and conformal calibration

CH. 5, 6

SLIDES 7–11

RQ4 · Limits

Will P, Q and V stay within limits?

One network for P, Q and V, and how useful it is

CH. 7

SLIDES 12–14

Back to Joanna

What does Joanna get, and what is left open?

Software, limits and next steps

SLIDES 15–18

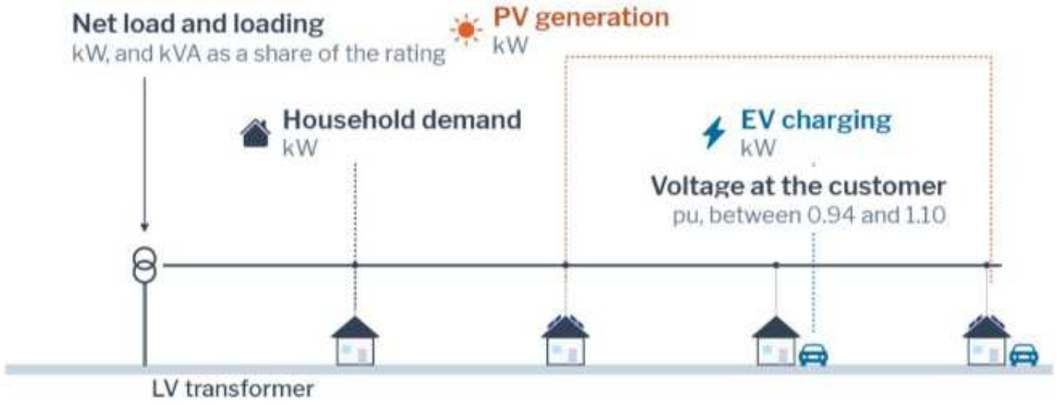
THESIS CLAIM

Each of the four questions tests one part of the claim: a lightweight network with calibrated uncertainty estimates can provide the forecasts required by Joanna's decisions.

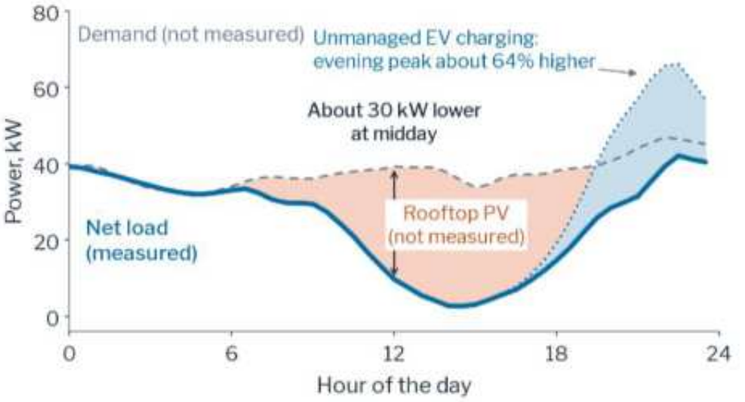
Forecasting net load at Joanna's substation is challenging

THE PROBLEM AND THE GOAL

Joanna's feeder



A summer day at her transformer



Less predictable group behaviour. A low-voltage substation aggregates few customers, so demand is less smooth than at higher voltage, and behind-the-meter PV makes the net load volatile. — Accurate

Growing uncertainty. PV and EV charging add uncertainty to demand and generation, so she must know how uncertain tomorrow is. — Reliable

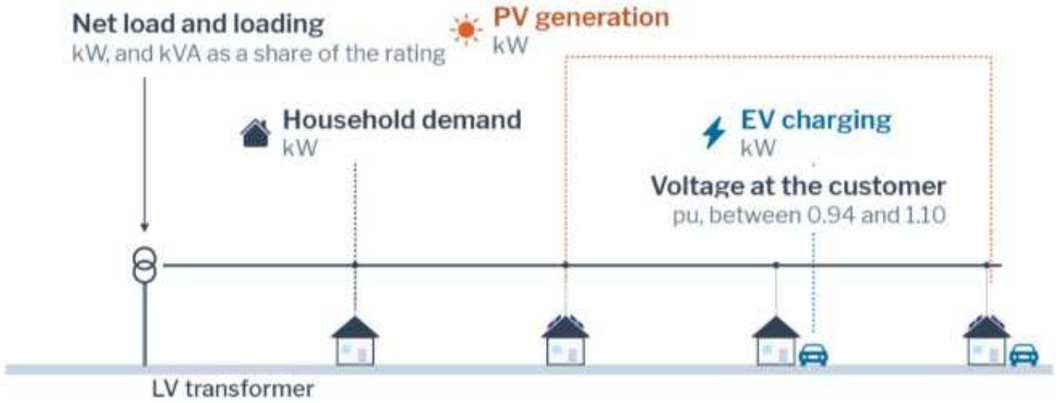
Scalability. A distribution operator runs thousands of such substations, so the forecast must be cheap to train and run at each of them. — Scalable

Illustration from real days of MLVS-PT: the mean of the four clearest and the four cloudiest summer weekdays; demand is taken as the cloudy-day net load and PV as the difference; the EV curve applies the 64% average peak increase reported by Byström (2022).

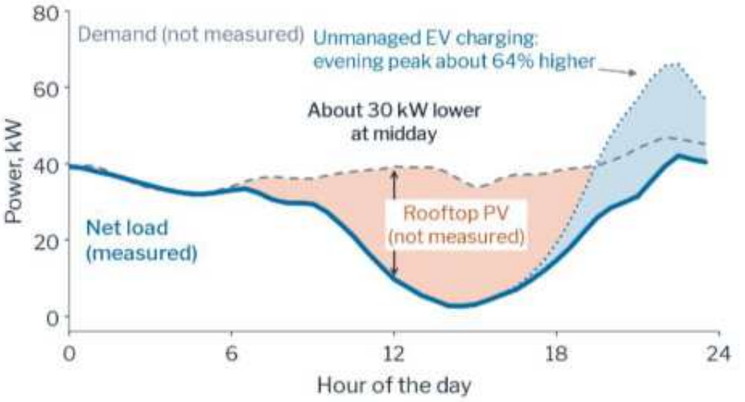
Forecasting net load at Joanna's substation is challenging

THE PROBLEM AND THE GOAL

Joanna's feeder



A summer day at her transformer



Less predictable group behaviour. A low-voltage substation aggregates few customers, so demand is less smooth than at higher voltage, and behind-the-meter PV makes the net load volatile. — Accurate

Growing uncertainty. PV and EV charging add uncertainty to demand and generation, so she must know how uncertain tomorrow is. — Reliable

Scalability. A distribution operator runs thousands of such substations, so the forecast must be cheap to train and run at each of them. — Scalable

GOAL

Develop, evaluate and demonstrate an accurate, reliable and scalable probabilistic forecasting approach for LV substations, so that Joanna can shift EV charging, schedule storage, reconfigure feeders and control voltage.

Illustration from real days of MLVS-PT: the mean of the four clearest and the four cloudiest summer weekdays; demand is taken as the cloudy-day net load and PV as the difference; the EV curve applies the 64% average peak increase reported by Byström (2022).

RESEARCH QUESTION 1 · CHAPTER 3 · SLIDES 2–3

Can I trust how forecasting models are compared?

FORMAL QUESTION

How can machine-learning algorithms for net-load forecasting at secondary substations be developed and evaluated, considering the influence of meteorological variables on forecasting accuracy?

PUBLISHED WORK

The experiment on exogenous variables is reported in A. Faustine, N. J. Nunes and L. Pereira, “Efficiency through Simplicity”, IEEE Transactions on Power Systems, 2025; the protocol of this part is the one used in all the papers of the thesis

A five-step protocol for comparing models

RQ1 · CHAPTER 3 · CONTRIBUTION

Develop · choose the inputs and window

Evaluate · benchmark the models

1

Signal characterisation

Lookback and lags read from the autocorrelation peaks.

→ **Lookback 48, lags 96 and 144**

2

Rank-fusion selection

Six measures rank each input; a Borda count fuses the ranks.

→ **Irradiance ranks first**

3

Rolling-origin backtest

Ten expanding folds from 12 months; 4-month tests moved by 2 months.

→ **10 folds, not one split**

4

Robust and proper scores

MAE and NRMSE for points; CRPS and Winkler score for intervals.

→ **Absolute error, not MAPE**

5

Significance tests

Diebold–Mariano on paired losses; model confidence set.

→ **Almost every pair differs significantly**

IMPLICATION

Joanna can repeat these five steps on her own data with Twiga and verify a claimed improvement before adopting it.

RESEARCH QUESTION 2 · CHAPTERS 4 AND 6 · SLIDES 4–6

Can I afford an accurate model at every substation?

FORMAL QUESTION

How can a lightweight DNN be designed to improve both computational efficiency and forecasting performance in LV net-load forecasting?

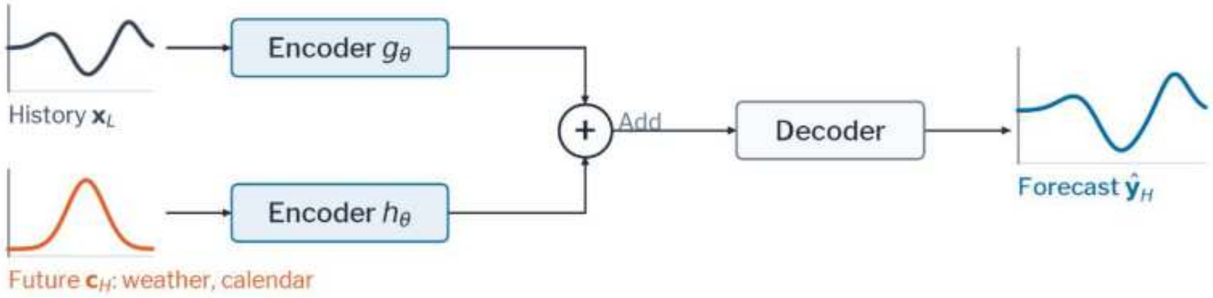
PUBLISHED WORK

A. Faustine, N. J. Nunes and L. Pereira, “Efficiency through Simplicity: MLP-based approach for net-load forecasting with uncertainty estimates in low-voltage distribution networks”, IEEE Transactions on Power Systems, 2025

MLPF and MLPGAM: two small encoders, one addition

RQ2 · CHAPTERS 4 AND 6 · MOTIVATION AND METHOD

MLPF · the design



$$\phi(g) = g_\theta(E(\mathbf{x}_L)), \qquad \phi(h) = h_\theta(E(\mathbf{c}_H))$$

$$\hat{\mathbf{y}}_H = \text{Decoder}(\phi(g) + \phi(h))$$

$$\mathcal{L} = \frac{1}{H}\sum_{t=1}^{H}\left[\lambda\,(y_t - \hat{y}_t)^2 + (1-\lambda)\,|y_t - \hat{y}_t|\right]$$

Adding beats attention. Replacing the attention layer by an addition lowers the MAE by 17.7% (MLVS-PT) and 20% (SPS-UK).

Foundation models. Pre-trained on many kinds of series, with little evidence yet on LV net load with PV.

Transformers. More than 300 times slower to infer than MLPF, and not consistently more accurate.

DLinear, N-BEATS, N-HiTS. Lighter, but designed around the history of the series, with covariates as an add-on.

MLPF. A small network for each substation, with an encoder of its own for the known future, the irradiance forecast and the calendar. The decoder accepts any head: point, quantile, Laplace, Bayesian or FPQR.

IMPLICATION

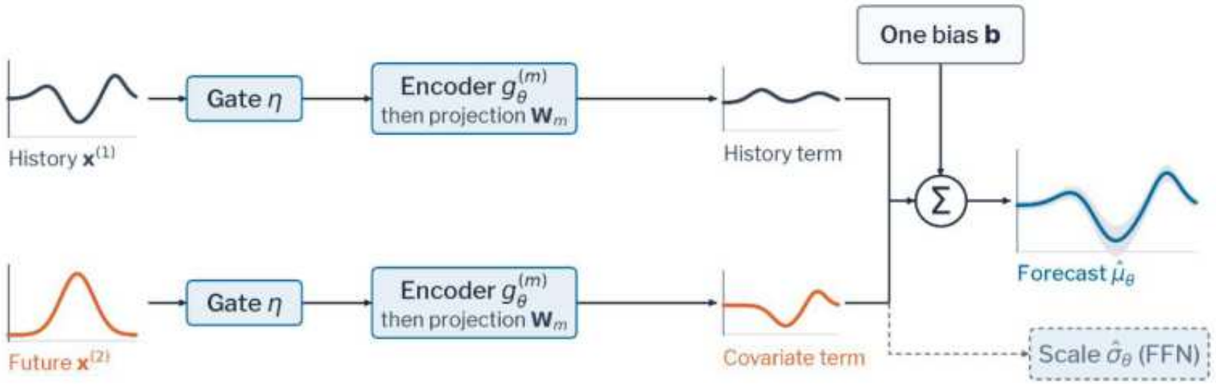
Joanna gets the accuracy of a deep model from a network small enough to run at every substation.

Schematic diagrams; the curves illustrate the inputs, outputs and terms and are not results.

MLPF and MLPGAM: two small encoders, one addition

RQ2 · CHAPTERS 4 AND 6 · MOTIVATION AND METHOD

MLPGAM · the same network with separate terms



$$\hat{\boldsymbol{\mu}} = \psi\left(\mathbf{b} + \sum_{m=1}^{M} \mathbf{W}_m^{\top}\, g_\theta^{(m)}\left(\mathbf{x}^{(m)} \odot \eta^{(m)}\right)\right), \qquad \eta^{(m)} = \sigma\left(\tau \pi^{(m)}\right)$$

One gate, encoder and bias-free projection for each input group; the terms are added to one shared bias.

⊕ **Addition already helped.** In MLPF, adding the two encoder outputs beat attention (MAE 17.7% and 20% lower), so addition is enough.

Why add every group. A sum of one term per input group splits each forecast exactly into a history term and a covariate term, so the contribution of each group can be read off.

Why a bias. One shared bias carries the level and the projections have none, so each term is a deviation from that level.

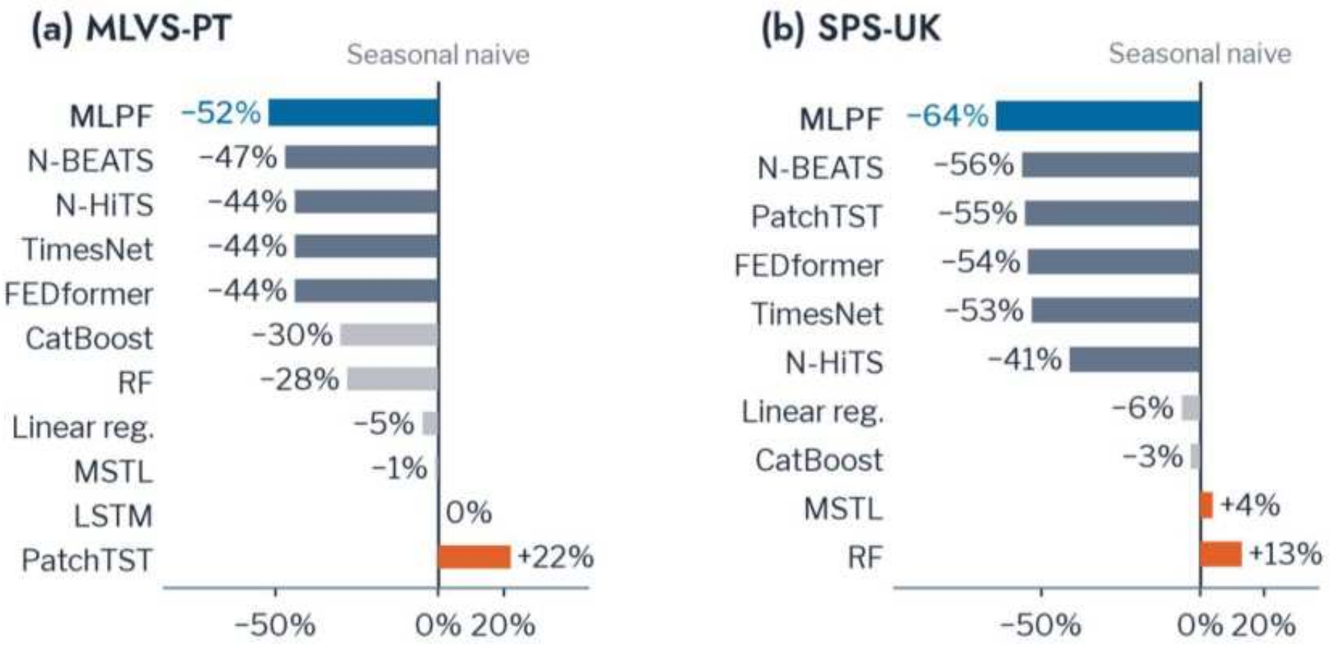
IMPLICATION

Joanna can read how much of each forecast comes from the weather forecast.

Schematic diagrams; the curves illustrate the inputs, outputs and terms and are not results.

MLPF is significantly the most accurate model

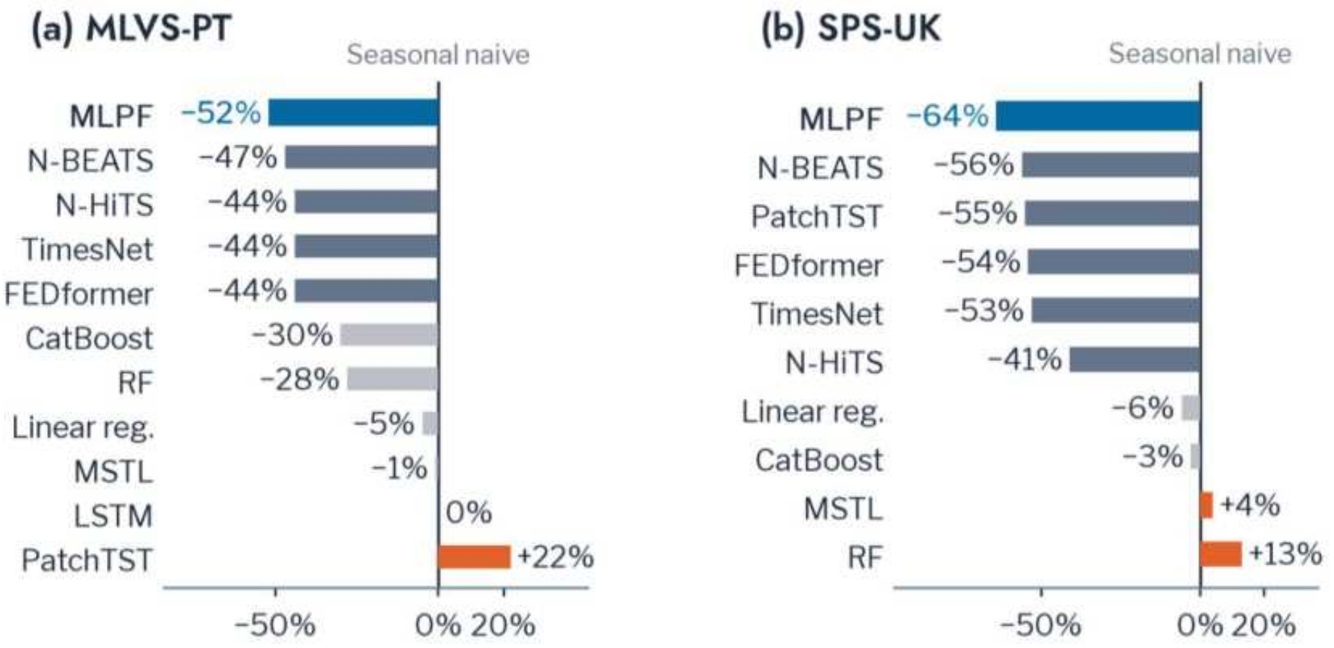
RQ2 · CHAPTER 4 · CONTRIBUTION



Change of the mean absolute error against the seasonal naive forecast over the backtesting folds: twelve models on MLVS-PT, eleven on SPS-UK. MLPF blue, deep models slate, tree and statistical models grey, worse than the naive forecast orange

MLPF is significantly the most accurate model

RQ2 · CHAPTER 4 · CONTRIBUTION



Change of the mean absolute error against the seasonal naive forecast over the backtesting folds: twelve models on MLVS-PT, eleven on SPS-UK. MLPF blue, deep models slate, tree and statistical models grey, worse than the naive forecast orange

52% and 64% lower MAE than the seasonal naive forecast, the largest reduction of any model

Significantly better. The Diebold–Mariano test finds MLPF better than every other model on both datasets, all tuned with the same budget; N-BEATS ranks second

About 12 ms to forecast six months, against 15 to 16 s for N-BEATS and N-HiTS and more than 300 times longer for PatchTST and FEDformer

IMPLICATION

Joanna can run one MLPF at every substation without giving up accuracy.

Thesis Chapter 4: change of the mean absolute error over the backtesting folds; all models tuned with the same Optuna budget; inference time for the six-month test window.

RESEARCH QUESTION 3 · CHAPTERS 5 AND 6 · SLIDES 7–11

How uncertain is tomorrow, and does a 90% interval achieve 90% coverage?

FORMAL QUESTION

How can lightweight DNN models be developed to produce well-calibrated probabilistic forecasts without compromising forecasting accuracy?

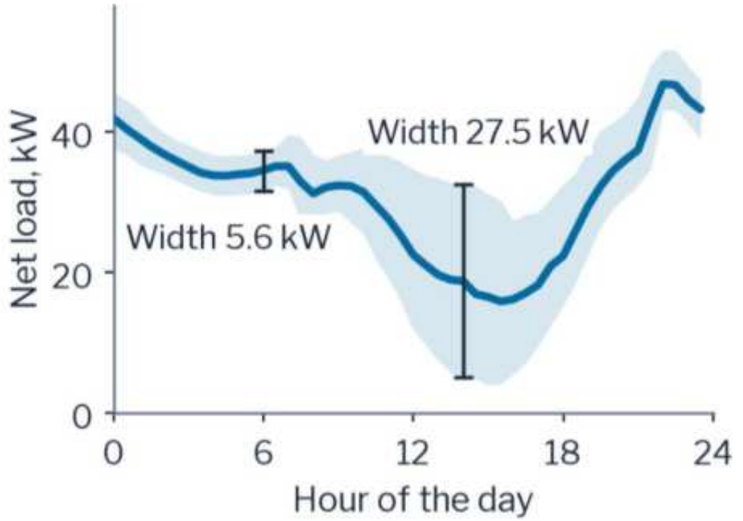
PUBLISHED WORK

A. Faustine and L. Pereira, “FPSeq2Q: fully parameterised sequence to quantile regression for net-load forecasting with uncertainty estimates”, IEEE Transactions on Smart Grid, 2022 · “Conformal multilayer perceptron-based probabilistic net-load forecasting for low-voltage distribution systems with PV generation”, IEEE SmartGridComm, 2024

Uncertainty is about five times larger at midday than at dawn

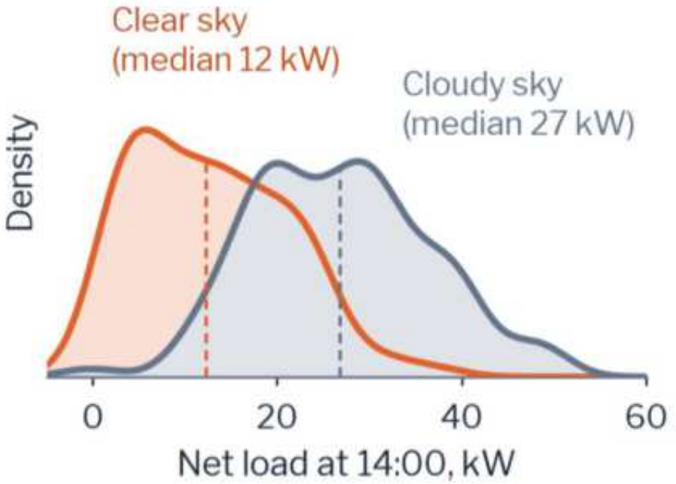
RQ3 · CHAPTER 5 · MOTIVATION

The interval width changes with the hour



Almost five times wider. In summer the 10 to 90% range of her net load is 5.6 kW wide at 06:00 and 27.5 kW wide at 14:00, when PV and passing cloud decide the outcome.

At midday there are two distinct outcomes



12 kW or 27 kW. At 14:00 a clear sky leaves a median net load of 12 kW and a cloudy sky 27 kW, so a single value cannot describe the day.

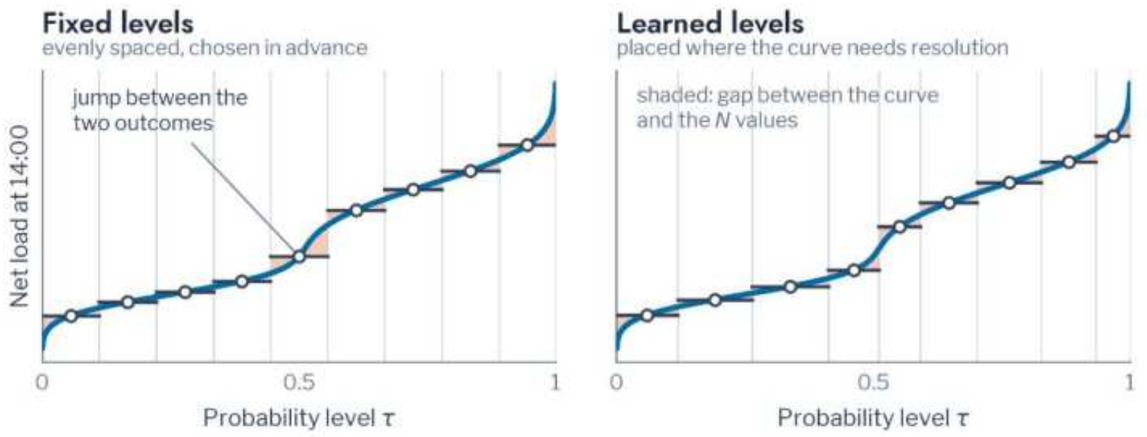
IMPLICATION

Joanna needs an interval that is narrow at dawn and wide at midday, so the quantile levels should be learned, not evenly spaced.

MLVS-PT, June to August, 92 days; clear and cloudy sky are the upper and lower third of the irradiance at 14:00.

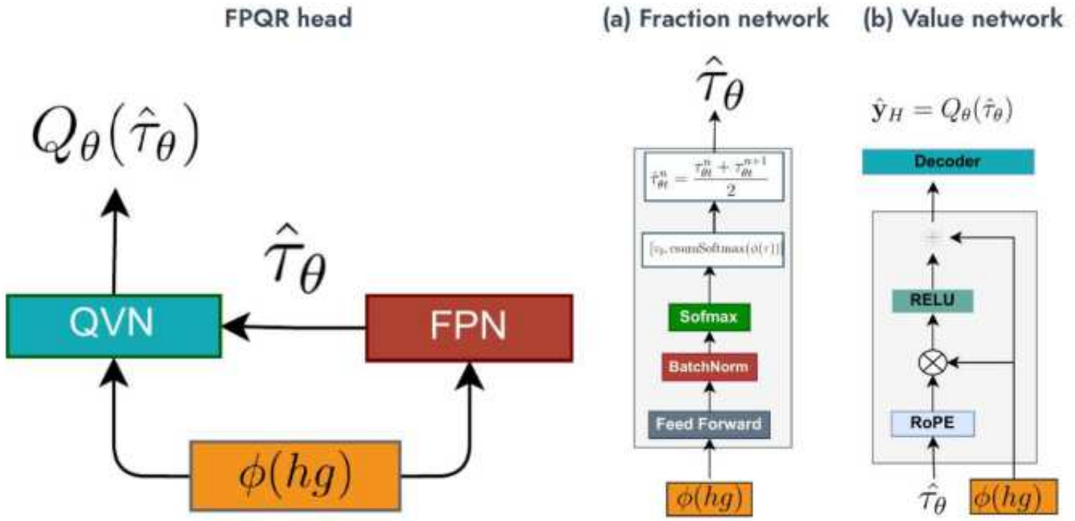
FPQR narrows the intervals, but its coverage is not controlled

RQ3 · CHAPTER 5 · CONTRIBUTION



Learn the levels, do not fix them

- Evenly spaced levels spend the same resolution at night, when net load is narrow with one mode, and at midday under PV, when it is broad with two.
- A fraction network places the levels and a value network gives the net load at each.
- No distribution is assumed.



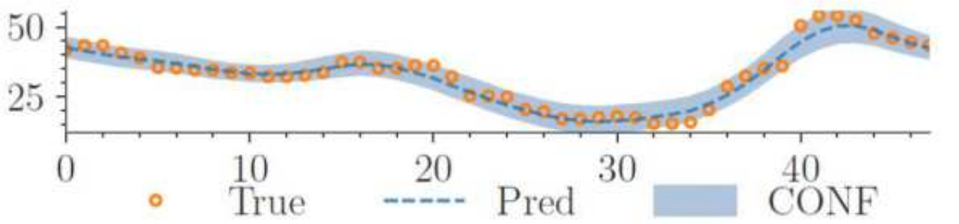
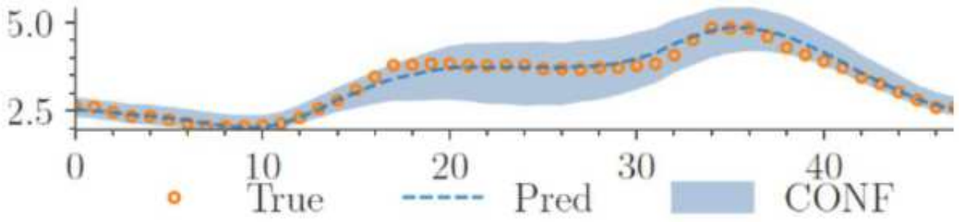
$$\mathcal{L}_\tau = \frac{1}{T}\sum_t \sum_n \rho^{\kappa}_{\hat{\tau}_n}(y_t - Q_\theta(\hat{\tau}_n)) \quad \text{(value network), with the non-crossing penalty } \mathcal{L}_{\text{nc}}$$

$$W_1 = \sum_n \int_{\tau^n}^{\tau^{n+1}} |Q(\tau) - Q(\hat{\tau}_n)| d\tau, \quad \mathcal{L}_{\mathcal{H}} = -\lambda_{\mathcal{H}} H(\Delta\tau) \quad \text{(fraction network)}$$

$$J = \mathcal{L}_\tau + \mathcal{L}_{\text{nc}} + W_1 + \mathcal{L}_{\mathcal{H}}, \quad \text{stop-gradient so that } \mathcal{L}_\tau \text{ cannot move the levels}$$

FPQR narrows the intervals, but its coverage is not controlled

RQ3 · CHAPTER 5 · CONTRIBUTION

Dataset	Model	NRMSE	PICP	NMPI	NWS	NCRPS	Example forecast
MLVS-PT	MLP-BNN	0.10	0.94	0.42	0.42	0.06	
	MLP-MCD	0.12	0.93	0.42	0.42	0.06	
	MLPQF	**0.09**	**0.78**	**0.21**	**0.25**	**0.04**	
SPS-UK	MLP-BNN	0.15	0.95	0.70	0.70	0.08	
	MLP-MCD	0.21	0.92	0.65	0.65	0.08	
	MLPQF	**0.13**	**0.84**	**0.23**	**0.28**	**0.05**	

Lower error, narrower intervals. MLPQF has the lowest NRMSE, NMPI and NCRPS on both datasets; the Bayesian and MC-dropout intervals are two to three times wider

Coverage below the target. PICP of 0.78 to 0.84 for MLPQF against 0.92 to 0.95, because the competing intervals are wide; coverage is not controlled

IMPLICATION

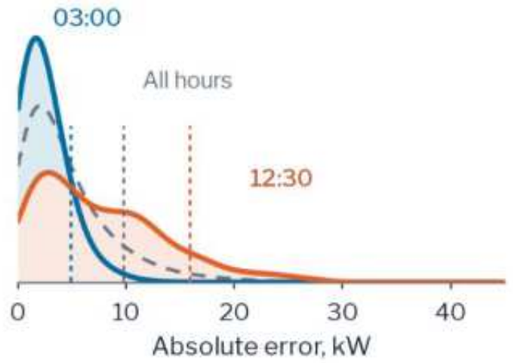
Joanna can add FPQR to her network for narrower intervals, but must calibrate its coverage before sizing storage.

Thesis Chapter 5; original paper (IEEE Transactions on Power Systems, 2024), two datasets.

Conformal calibration sets the interval width from held-out errors

RQ3 · CHAPTER 6 · METHOD

Why · the error depends on the hour



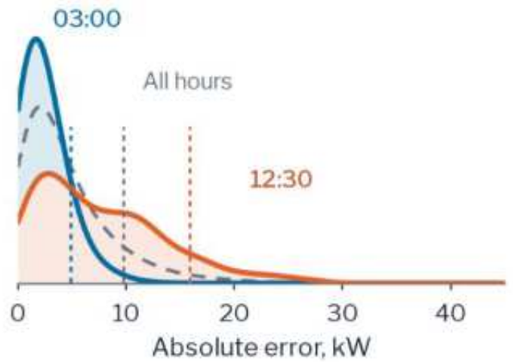
Absolute error at 03:00, at 12:30 and over all hours, MLVS-PT; dotted lines mark the 90% quantile of each

One pooled width suits none of them. Too wide at dawn and too narrow at midday, so coverage changes with the hour

Conformal calibration sets the interval width from held-out errors

RQ3 · CHAPTER 6 · METHOD

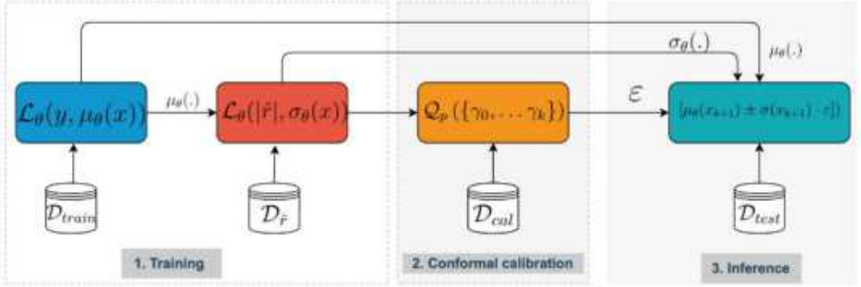
Why · the error depends on the hour



Absolute error at 03:00, at 12:30 and over all hours, MLVS-PT; dotted lines mark the 90% quantile of each

One pooled width suits none of them. Too wide at dawn and too narrow at midday, so coverage changes with the hour

How · width from held-out errors



$$\gamma_{k,h} = \frac{|y_{k,h} - \hat{\mu}_h(\mathbf{x}_k)|}{\hat{\sigma}_h(\mathbf{x}_k)}, \qquad \varepsilon_h = Q_{1-\alpha}(\gamma_{1,h}, \ldots, \gamma_{N,h})$$

$$\mathcal{C}_h(\mathbf{x}) = [\hat{\mu}_h(\mathbf{x}) \pm \varepsilon_h \hat{\sigma}_h(\mathbf{x})]$$

Adaptive conformal calibration, thesis Chapter 6: training, calibration and inference on separate data

One width per lead time. The 90% quantile of the scaled absolute error, from held-out days; no error distribution is assumed

IMPLICATION

Joanna keeps one forecast and sets the 90% width from its own recent errors, not from an assumed distribution.

Thesis Chapter 6, MLVS-PT; conformal calibration with the absolute score.

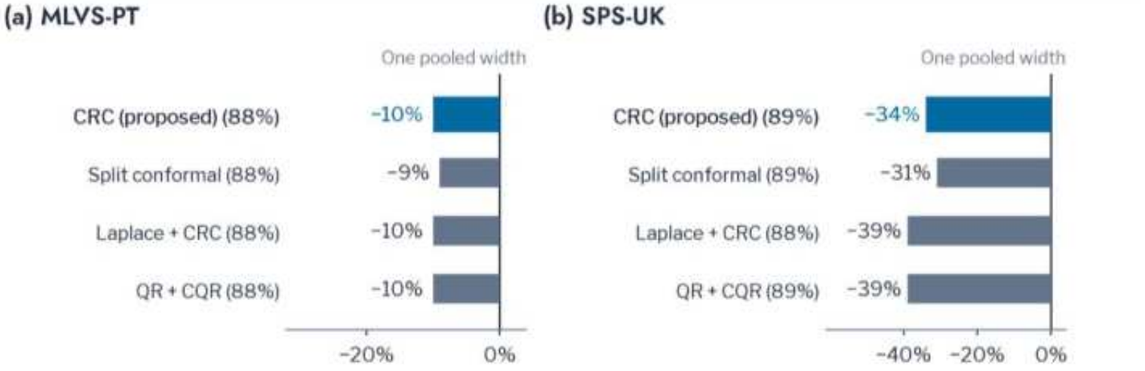
Calibration brings QR and Laplace to nominal coverage

RQ3 · CHAPTER 6 · CONTRIBUTION

Raw to calibrated coverage



Winkler score against one pooled width



Coverage moves to nominal. QR from 0.71 and 0.83 to 0.88 and 0.89; Laplace from 0.79 and 0.86 to 0.88, with the point forecast unchanged

CRC and split conformal. Winkler score 9% to 10% lower on MLVS-PT and 31% to 34% lower on SPS-UK than one pooled width

Any head can use it. Calibrated Laplace and QR match CRC on MLVS-PT and are 5 points lower on SPS-UK, whose mean forecast is more accurate

IMPLICATION

Joanna can calibrate the same lightweight point forecast with CRC and its scale, or recalibrate a Laplace or quantile head on top of MLPF.

Daily Winkler score on the distinct test days (about 395 per dataset).

RESEARCH QUESTION 4 · CHAPTER 7 · SLIDES 12–14

Can a lightweight network forecast P, Q and V together, and how useful is the forecast?

FORMAL QUESTION

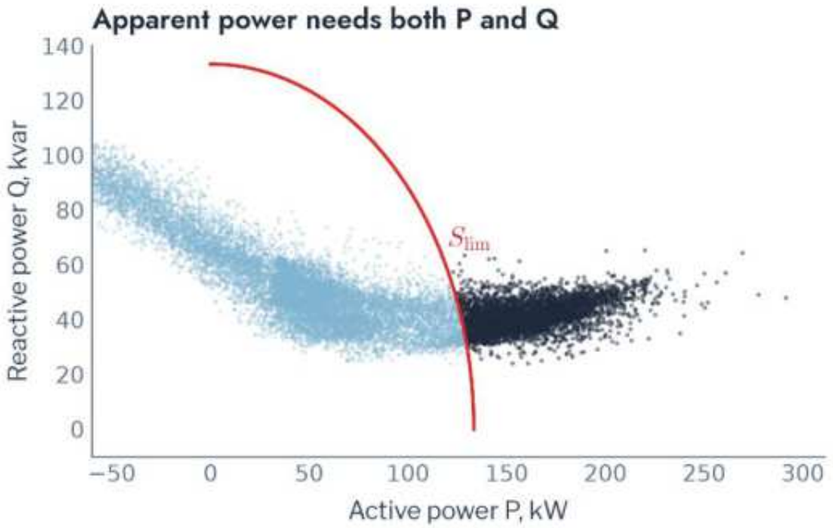
How can lightweight DNN-based models be designed to generate well-calibrated multi-target probabilistic forecasts of LV network variables for uncertainty-aware network decision-making?

PUBLISHED WORK

A. Faustine and L. Pereira, probabilistic forecasting of voltage, reactive and active power for the resilience of low-voltage systems, CIRED Chicago Workshop, 2024

Can one lightweight network forecast P, Q and V together?

RQ4 · CHAPTER 7 · FORECASTING P, Q AND V



Simulated weak feeder (DSS-UK), 40% EV and 40% PV, one year of half-hours. The red arc is the apparent-power limit and the dark points lie beyond it

QUESTION

Can the lightweight MLPF forecast P, Q and V together, with reliable intervals?

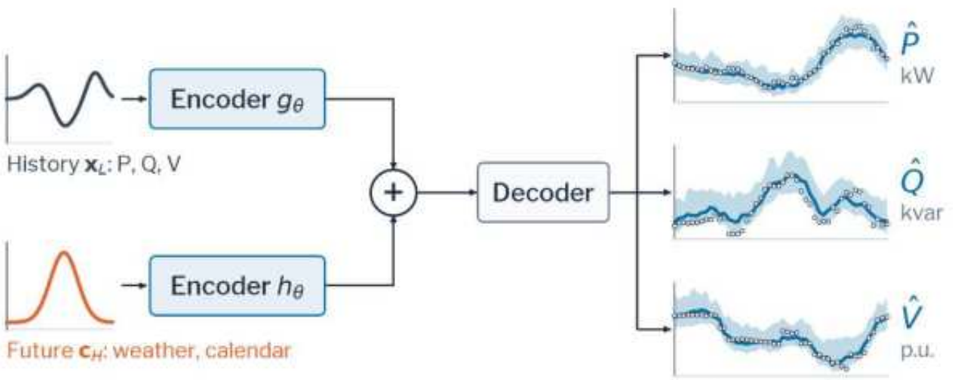
Limits involve several variables

- **Transformer loading** is apparent power, from P and Q.
- **Customer voltage** must stay above 0.9 p.u.
- **A decision needs all three** at once.

One network, not three

- **P, Q and V share** weather, behaviour and PV and EV growth.
- **One network** shares what it learns and costs less.

Can one lightweight network forecast P, Q and V together?



The MLPF of the previous part with three outputs, shown for one test day of the simulated weak feeder: forecast, 90% interval and observed values

$$\mathbf{y}_t = [P_t,\ Q_t,\ V_t]^\top, \qquad \Phi = g_\theta(\mathbf{x}_L, \mathbf{c}_H)$$

$$\hat{F}_{d,h}(y) = F_{d,h}(y \mid \mathbf{x}_L, \mathbf{c}_H), \qquad d = 1, 2, 3$$

$$\mathcal{L}_{\mathrm{MT}} = \frac{1}{DH} \sum_{d=1}^{D} \sum_{h=1}^{H} \ell_{\mathrm{head}}(y_{d,h})$$

A head on the shared representation gives the distribution of each target and lead time; each head has its own loss ℓ_{head}, which the general loss averages over targets and lead times

QR and FPQR

$$\ell_{\mathrm{head}} = \sum_{\tau} \rho_\tau(y - Q(\tau))$$

pinball loss over the levels

Laplace

$$\ell_{\mathrm{head}} = \frac{|y - \mu|}{b} + \log(2b)$$

negative log-likelihood

BNN

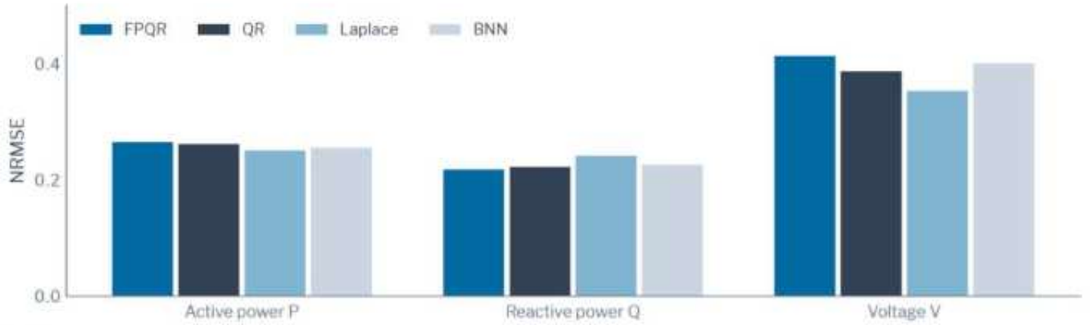
$$\ell_{\mathrm{head}} = -\mathrm{ELBO}$$

evidence lower bound

Can one lightweight network forecast P, Q and V together?

RQ4 · CHAPTER 7 · FORECASTING P, Q AND V

Experiment 1 · four heads, four records



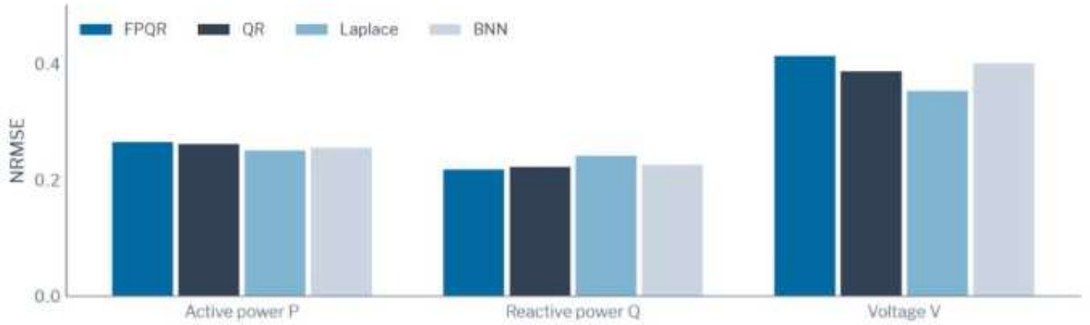
NRMSE of each target, measured substation

- **Voltage is the hardest target.** NRMSE of 0.35 to 0.41 against 0.22 to 0.27 for P and Q; QR and Laplace beat FPQR on V
- **The heads are otherwise alike.** Average ranks are 2.46 to 2.57 on the measured substation

Can one lightweight network forecast P, Q and V together?

RQ4 · CHAPTER 7 · FORECASTING P, Q AND V

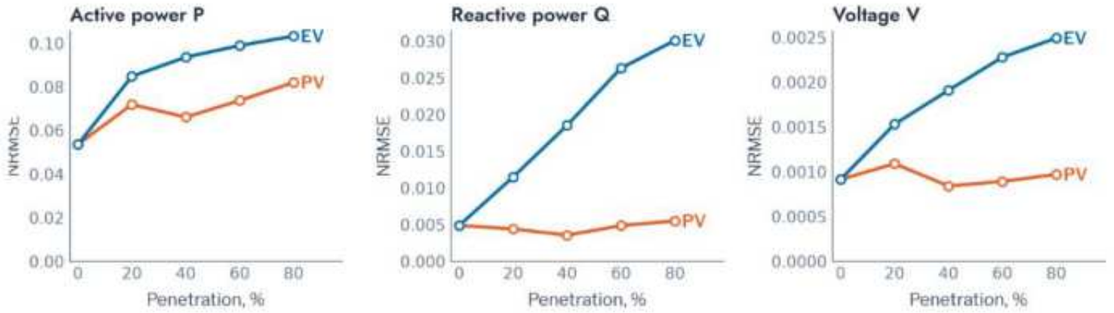
Experiment 1 · four heads, four records



NRMSE of each target, measured substation

- **Voltage is the hardest target.** NRMSE of 0.35 to 0.41 against 0.22 to 0.27 for P and Q; QR and Laplace beat FPQR on V
- **The heads are otherwise alike.** Average ranks are 2.46 to 2.57 on the measured substation

Experiment 2 · PV and EV penetration, DSS-EU



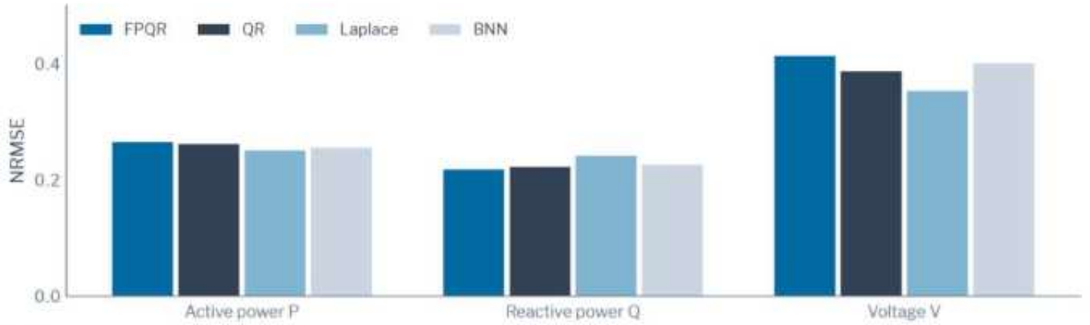
NRMSE of P, Q and V as PV or EV penetration grows

- **EV matters more than PV.** At 80% EV the error of P is about twice its no-DER value and that of Q about six times; PV changes Q and V little
- **Q and V are the most sensitive.** They deteriorate steadily with EV charging across the whole range

Can one lightweight network forecast P, Q and V together?

RQ4 · CHAPTER 7 · FORECASTING P, Q AND V

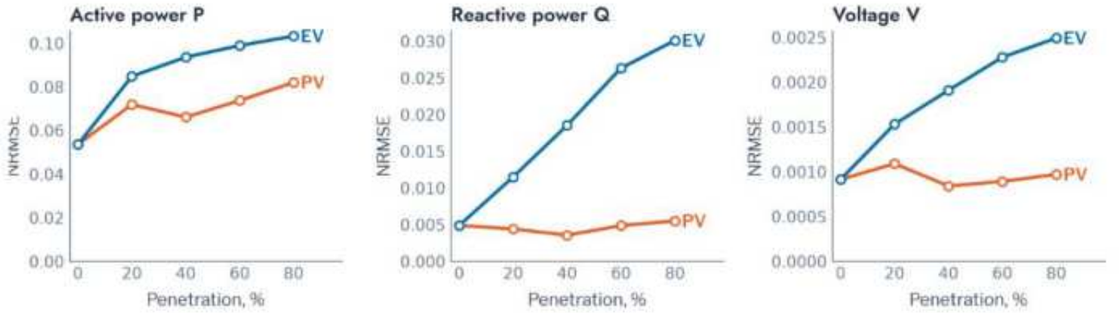
Experiment 1 · four heads, four records



NRMSE of each target, measured substation

- **Voltage is the hardest target.** NRMSE of 0.35 to 0.41 against 0.22 to 0.27 for P and Q; QR and Laplace beat FPQR on V
- **The heads are otherwise alike.** Average ranks are 2.46 to 2.57 on the measured substation

Experiment 2 · PV and EV penetration, DSS-EU



NRMSE of P, Q and V as PV or EV penetration grows

- **EV matters more than PV.** At 80% EV the error of P is about twice its no-DER value and that of Q about six times; PV changes Q and V little
- **Q and V are the most sensitive.** They deteriorate steadily with EV charging across the whole range

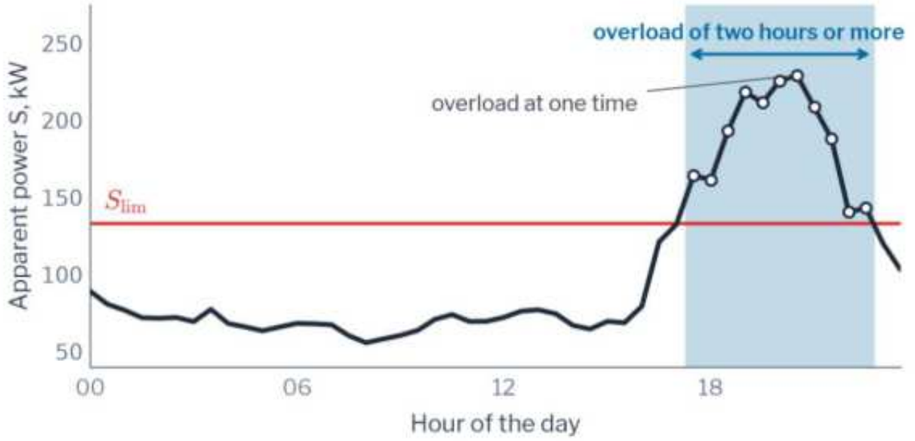
IMPLICATION

Joanna needs one network for P, Q and V, with the dependence between their intervals added.

Thesis Chapter 7, Experiments 1 and 2; forecast example on the simulated DSS-UK weak feeder.

From forecasts to the probability of an event

RQ4 · CHAPTER 7 · HOW USEFUL IS THE FORECAST?



One day of the simulated weak feeder (DSS-UK): apparent power against its limit, the 98th percentile with no DER

$$S_t = \sqrt{P_t^2 + Q_t^2}, \qquad S_{\text{lim}} = 98\text{th percentile of } S \text{ with no DER}$$

$$\text{one variable}: \quad p = \Pr(Y > y^*) = 1 - F(y^*)$$

$$\text{one half-hour}: \quad p = \frac{1}{M}\sum_{m=1}^{M} \mathbf{1}\left[S_t^{(m)} > S_{\text{lim}}\right], \quad y^{(m)} = F^{-1}(u^{(m)})$$

$$\text{two hours}: \quad p = \frac{1}{M}\sum_{m=1}^{M} \mathbf{1}\left[S_t^{(m)} > S_{\text{lim}}, \ldots, S_{t+3}^{(m)} > S_{\text{lim}}\right]$$

The heads give marginals

Quantile and Laplace heads give one distribution for each variable and half-hour, and their intervals cover P, Q and V together in only 29% to 81% of cases, so the scenarios need the distribution function and a dependence

From forecasts to the probability of an event

1 **The distribution function gives the draws**

Each curve is one draw: a random level u is passed through the inverse of F.

$$\tilde{y}^{(m)} = F^{-1}(u^{(m)}), \quad u^{(m)} \sim \mathcal{U}(0,1)$$

2 **Reordering keeps the values and fixes the pairing**

The same draws are assigned to the rank pattern of recent day m, which is the empirical copula.

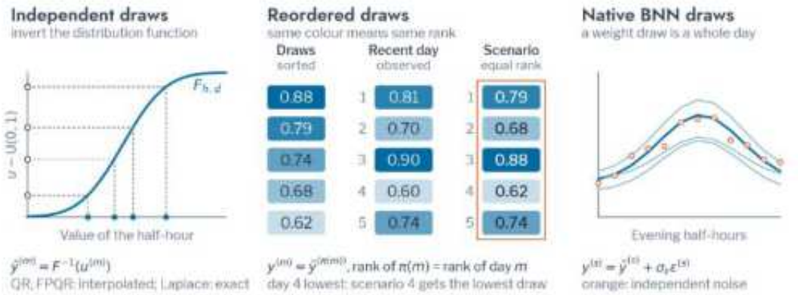
$$y_h^{(m)} = \tilde{y}_h^{(\pi_h(m))}, \quad \text{rank } \pi_h(m) = \text{rank of recent day } m$$

3 **The probability is the share of curves that hit the event**

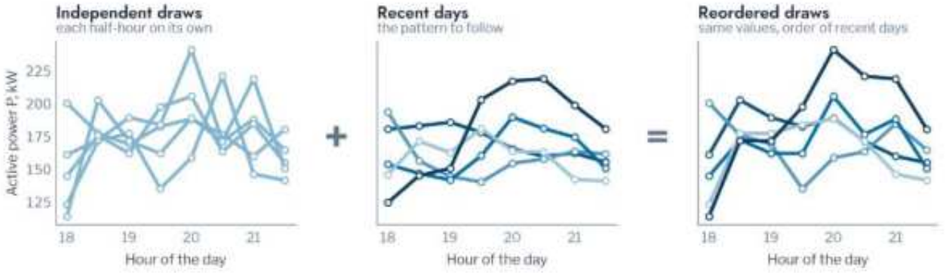
Count the scenarios that exceed the limit for two hours or more; for a single half-hour the count reduces to 1 − F.

$$p = \frac{1}{M}\sum_{m=1}^{M} \mathbf{1}[\text{scenario } m \text{ has the event}]$$

$$\text{one half-hour}: \quad p = 1 - F(y^*)$$



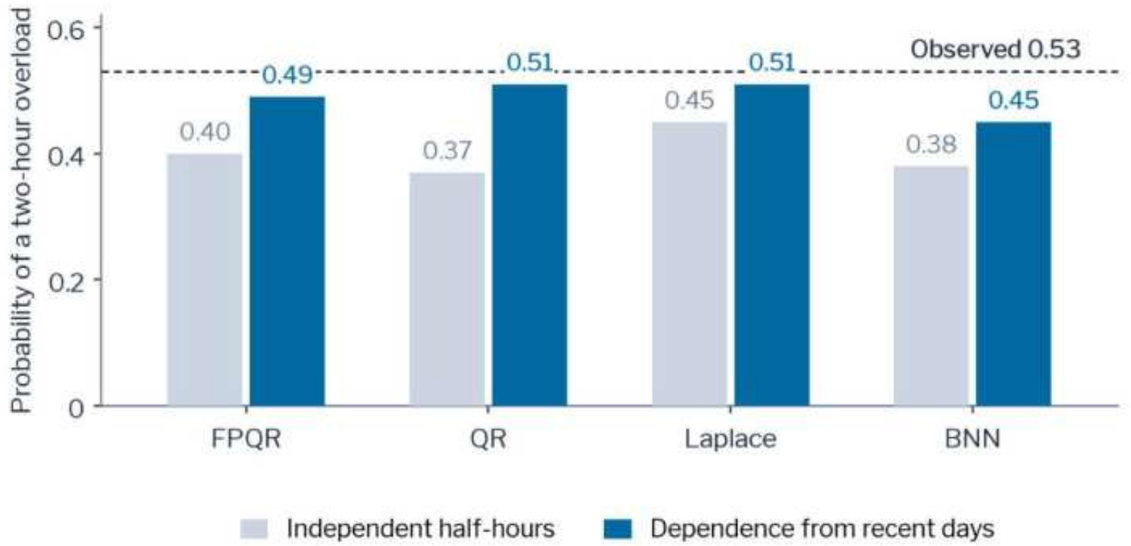
The Schaake shuffle: draws are sorted, then reassigned to follow the ranks of a recent day



Real data: five scenarios of active power over four evening hours from the Laplace head of the thesis run (27 September 2019), and the five recent days before it that they are reordered to follow

From forecasts to the probability of an event

RQ4 · CHAPTER 7 · HOW USEFUL IS THE FORECAST?



Forecast probability that an overload lasts at least two hours, on the UK feeders, against the observed rate of 0.53. Blue: the same forecasts reordered to follow recent days, without retraining

Independent half-hours: too low. Treating each half-hour on its own gives 0.37 to 0.45, against 0.53 observed

Reordered scenarios: close. Following recent days raises the estimate to 0.49 to 0.51 and lowers the Brier score for every head, without retraining

IMPLICATION

Joanna can obtain the probability of a persistent event from the same forecasts, without retraining.

Thesis Chapter 7, DSS-UK feeders: two-hour overload above the pre-DER 98th-percentile apparent power; dependence by the Schaake shuffle.

BACK TO JOANNA · CONCLUSION · SLIDES 15–18

What does Joanna get, and what is left open?

CLAIM OF THE THESIS

A lightweight network with calibrated uncertainty estimates can provide the forecasts required by her four decisions.

The four decisions: assessment

BACK TO JOANNA · CONCLUSION

Shift EV charging



Tomorrow's peak, from a verified and affordable model

−52% and −64%

mean absolute error against the seasonal naive forecast; 12 ms for six months

Supported RQ1 · RQ2

Schedule storage or flexibility



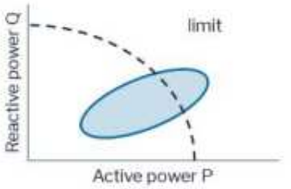
How uncertain tomorrow is, with nominal interval coverage

88% to 89%

coverage of the 90% interval after calibration; raw QR and Laplace 71% to 86%

Supported RQ3

Reconfigure feeders



The risk of an overload, from P and Q together

0.45 to 0.51

forecast chance of a two-hour overload with the dependence added, against 0.53 observed

Partly supported RQ4

Control voltage



The risk of a voltage violation

0.35 to 0.41

NRMSE of voltage, the hardest of the three targets, against 0.22 to 0.27 for P and Q

Partly supported RQ4

IMPLICATION

Two of Joanna's four decisions are fully supported and two in part.

Schematic illustrations; figures from thesis Chapters 3 to 7.

Twiga implements the workflow

BACK TO JOANNA · SOFTWARE

1

Compare models

Select the inputs, backtest on rolling windows and test the differences.

`backtesting()`

→ **Protocol, Ch. 3**

2

Forecast

Train the lightweight MLPF or MLPGAM on her own substation record.

`fit()`

→ **Ch. 4, 6**

3

Quantify uncertainty

Learn the quantile levels of the predictive distribution with FPQR.

`predict_interval()`

→ **Ch. 5**

4

Calibrate

Set the interval width for each lead time from held-out errors.

`calibrate()`

→ **Ch. 6**

5

Forecast P, Q and V

One network for P, Q and V, with the dependence between them added for events that last.

`schaake_shuffle()`

→ **Ch. 7**

IMPLICATION

Joanna can run the same open-source workflow on the data of every substation.

Twiga: pypi.org/project/twiga (method names as in the library); Popo: pypi.org/project/popo-flow, which built the simulated feeders of Chapter 7.

Open questions

CONCLUSION · LIMITS AND NEXT STEPS

Measured MLVS-PT SPS-UK Simulated DSS-EU DSS-UK DSS-US

Two PV-dominated sites

Not claimed for wind-, heating- or EV-dominated networks.

NEXT STEP
Other resource mixes.

Simulated P, Q and V records

Three of four records are simulated, so decision values are indicative.

NEXT STEP
Measured P, Q and V.

Coverage across seasons

Coverage weakens after the winter transition.

NEXT STEP
Seasonal and online calibration.

Weather-forecast uncertainty

The irradiance forecast error is not propagated.

NEXT STEP
Propagate it.

Rare overloads

The copula from recent days gives few rare events.

NEXT STEP
Condition it on weather or season.

IMPLICATION
Joanna's decisions are supported today; extending them needs other networks, measured multi-target records and seasonal calibration.

Thank you

Questions are welcome



github.com/sambaiga/twiga-thesis